

(19) Europäisches Patentamt  
European Patent Office  
Office européen des brevets



(11) EP 1 051 001 A2

(12)

## EUROPEAN PATENT APPLICATION

(43) Date of publication:  
08.11.2000 Bulletin 2000/45

(51) Int Cl.<sup>7</sup>: H04L 12/56, H04Q 11/04

(21) Application number: 00111248.1

(22) Date of filing: 28.01.1994

(84) Designated Contracting States:  
AT BE CH DE DK ES FR GB GR IE IT LI LU MC NL  
PT SE

(30) Priority: 29.06.1993 US 11378

(62) Document number(s) of the earlier application(s) in  
accordance with Art. 76 EPC:  
94905904.2 / 0 681 770

(71) Applicant: TELEFONAKTIEBOLAGET LM  
ERICSSON  
126 25 Stockholm (SE)

(72) Inventor: Petersen, Lars-Göran  
147 42 Tumba (SE)

(74) Representative: Wennerholm, Kristian  
Ericsson Radio Systems AB  
Common Patent Department  
164 80 Stockholm (SE)

Remarks:

This application was filed on 25 - 05 - 2000 as a  
divisional application to the application mentioned  
under INID code 62.

(54) Controlled access ATM switch

(57) An asynchronous transfer mode (ATM) switch  
in which access to a switchcore matrix is monitored and  
controlled through the logic and buffering functions of  
switchports connected thereto. The switchcore is greatly  
simplified by moving the logic and buffering functions to the  
switchports. The switchcore matrix comprises a  
plurality of rows, columns, and crosspoints thereof, pro-

viding routing paths for the routing of information cells  
from input points to output points on the matrix. Single-  
store buffers in the switchcore matrix enable temporary  
storage and hand-off of individual information cells as  
they pass through the matrix. The simplicity of the  
switchcore matrix enables it to be constructed on a sin-  
gle integrated circuit.




FIG. 3

**Description****BACKGROUND OF THE INVENTION****5 Field of the Invention**

**[0001]** This invention relates to Asynchronous Transfer Mode (ATM) switches, and in particular, to the implementation of flow control and isochrone traffic within ATM switching systems.

**10 Description of Related Art**

**[0002]** Asynchronous Transfer Mode (ATM), also known as "cell relay", is a telecommunications standards-based technology designed to meet the demand for the public network to simultaneously multiplex and switch data over a wide spectrum of data rates. These requirements are the result of emerging multimedia, high-speed data and imaging applications. ATM is a statistical multiplexing and switching method which is based on fast packet switching concepts, and is a radical departure from the circuit switching techniques that are used by today's digital switches. ATM provides dedicated circuits for voice, data, and video communications by dividing the information flow within each of these three types of traffic into individual "cells", each cell including an address or directions specifying the location to which the information carried within the cell should be delivered. Direction instructions are added to the information carried by the cell in the form of a label, which is processed by the ATM switch as the cell is routed through the switch.

**[0003]** Several factors drive the design of broad-band ATM switching architectures:

1. The need to accommodate a wide range of traffic types from voice to video to data;
2. The high speed at which the switch has to operate (from 155 Mb/s to over 1.2 Gb/s); and
3. The burst-like nature of data communications.

**[0004]** If communications networks continue to be deployed with large switches in central locations, then a large-scale ATM switch will be necessary. If such a switch is to serve 50,000 to 100,000 customers, each operating at the SONET STS-3 rate (155 Mb/s), then the total customer access capacity at the switch-customer interface is about 10 terabits per second (Tb/s) in each direction. If only one-in-ten customers use their assigned access capacity at any one time, then the core of this large-scale ATM switch must be capable of switching about 1 Tb/s of traffic, which is several orders of magnitude larger than the capacity of today's local digital switches.

**[0005]** Several high-performance packet switching fabrics have been proposed in the past. These switch fabrics can be categorized into different architectures - internal buffer, input buffer, output buffer, shared buffer, or various combinations of these. Internal-buffered switches include the buffered banyan network. With the assumption of having single-cell buffers at the intermediate stage, and a balanced and uniform traffic pattern, the banyan switch's maximum throughput is only about 45% of that required for large-scale ATM switches. Input-buffered architectures include Batcher-banyan networks with ring reservation, or a self-routing crossbar network with parallel, centralized contention resolution. Because of head-of-line (HOL) blocking, its maximum throughput is about 58% of that required. Certain techniques, such as allowing two cells of each input port to compete with others increases the maximum throughput of input-buffered architectures to approximately 70%.

**[0006]** The other types of ATM switch architectures each have their own advantages. Switches with output buffering, for example, have been proven to give the best delay/throughput performance in large-scale switch architectures. The shared buffer architectures have been shown to improve memory utilization significantly. Other switches in the prior art include those equipped with mixed input and output buffers, and a Sunshine switch implemented with both internal and output buffers. Besides point-to-point switches, several multicast ATM switches have also been proposed.

**[0007]** Each type of switch architecture has its own advantages and disadvantages. For example, the Batcher-banyan network has fewer switch elements than a crossbar network does, but it has more difficulty in synchronizing all signals in every stage because interconnection wires are not identical between stages, and the wire-length difference increases as the network grows. Conversely, the crossbar network has more uniform and regular interconnections, resulting in easier synchronization, but it has more switch elements.

**[0008]** All of the prior art switches, and most of the current research in the area of ATM switching, is oriented toward developing switchcores of greater magnitude and complexity in order to provide the switching capacity necessary for a large-scale central switch operating under its maximum projected traffic load. Networks utilizing a dozen or more ATM chips have been designed in such switches in order to provide the large buffers and multiple pathways necessary to ensure a high probability that a cell will pass through the switchcore. There is also a great need, however, for high

quality ATM switches which are optimally designed for smaller relay nodes within various communications networks. None of the prior art ATM switch architectures, large or small, solve the capacity, throughput, and loss problems using access control, and none are capable of providing isochronal service.

5      **SUMMARY OF THE INVENTION**

[0009] In one aspect, the present invention is an asynchronous transfer mode (ATM) switch comprising a switchcore matrix and a plurality of switchports electronically connected to the switchcore matrix at input and output points. The switchports transmit and receive information cells from the switchcore matrix. The switchcore matrix provides routing paths for the routing of the information cells from the input points to the output points on the matrix through a plurality of rows, columns, and crosspoints thereof. The switchcore has multicast and broadcast capability. The switchports provide the interface between the ATM switch and external communications devices. The switchports also interface with the switchcore matrix by means of a switchcore interface, and control access to the switchcore matrix by means of feedback information from the switchcore matrix crosspoints. Access to the switchcore matrix may be controlled by one or more input buffers which store information cells until selected routing paths in the switchcore matrix are free. A plurality of switchcore matrices may be link-coupled to enhance switch performance.

[0010] In another aspect, the present invention is a method for controlling the flow of information cells within a communications system. The method begins by providing selectable routing paths for the routing of information cells from input points to output points of a switchcore matrix having a plurality of rows, columns, and crosspoints thereof. A plurality of switchports are then electronically connected to the input and output points of the switchcore matrix to transmit information cells thereto and receive information cells therefrom. This is followed by connecting each of the switchports to an external information cell communications device, and controlling access to the switchcore matrix available to each of the information cells. The step of controlling access to the switchcore matrix may also include storing the information cells in one or more input buffers located within each of the switchports until selected routing paths in the switchcore matrix are free.

[0011] It is an object of the present invention to provide an ATM switch with access control, which enables the quality of the connection through the switch to be controlled from units connected thereto, and eliminates the need for large buffers in the switchcore.

[0012] It is another object of the present invention to provide an ATM switch which maximizes the use of available bandwidth for data communications traffic, which is burst-like in nature, and which greatly reduces the loss rate of the switch. Cell loss often occurs in prior art switches when one stage of a switch transmits an information cell when the buffer of the receiving stage is full. When utilizing access control, information cells are held in input buffers, which are sized for the type of communications to be handled, until output buffers or ports are available. Losses are greatly reduced because they only occur if the input buffers are overloaded, and if the input buffers are properly sized, overloading is very rare.

[0013] It is still another object of the present invention to provide an ATM switch with properties which are similar to a local area network (LAN), thereby enabling the switch to better handle future data communications demands in public networks.

[0014] It is still another object of the present invention to provide an ATM switch that accommodates communications devices of differing speeds, thereby enabling new devices and future upgrades with higher speed capabilities to be attached, e.g., SONET STS 12c devices, without affecting lower speed devices which are already attached. This objective assures upgrading of existing equipment without the need to replace the entire ATM switch, provided that the switch core is upgraded at the same time to handle the increased speed.

[0015] It is still yet another object of the present invention to provide an ATM switch which may be modified to provide a predetermined delay when transmitting information cells, thereby enabling so-called isochronal traffic.

**BRIEF DESCRIPTION OF THE DRAWING**

[0016] The invention will be better understood, and its numerous objects and advantages will become apparent to those skilled in the art by reference to the following drawing in which:

FIG. 1 is a simplified block diagram of the major components of a controlled access ATM switch of the type employed in the system of the present invention;

FIG. 2 is a simplified functional block diagram of the ATM switch of FIG. 1, illustrating the data flow between switchports and the switchcore, and showing where input buffers are placed in the switchports in the system of the present invention;

FIG. 3 is a simplified functional block diagram of the ATM switch illustrating the manner in which operation and maintenance functions are monitored and controlled from the switchports in the system of the present invention;

- FIG. 4 is a simplified block diagram illustrating how all communications devices connected to the controlled access ATM switch of the present invention have access to the switchcore via the switchports;
- 5 FIG. 5 is a simplified functional block diagram illustrating the logic switching and space switching functions performed by the controlled access ATM switch of the present invention;
- FIG. 6 is a simplified functional block diagram illustrating the logic and space switching functions when several controlled access ATM switches are link-coupled in a matrix architecture;
- 10 FIG. 7 is a simplified functional block diagram illustrating the structure and relationships between the three communications protocol levels which perform the logic and space switching functions in the controlled access ATM switch of the present invention.
- FIG. 8 is a block diagram illustrating the physical lines of the switchcore interface (SCI) between one switchport and one plane of the switchcore;
- 15 FIG. 8a is a byte map of a generic cell as it is sent in each direction over the switchcore interface (SCI);
- FIG. 8b is a byte map of an information (traffic) cell as it is sent in each direction over the switchcore interface (SCI);
- FIG. 8c is a byte map of a maintenance cell as it is sent in each direction over the switchcore interface (SCI);
- 20 FIG. 8d is a byte map of an idle cell as it is sent in each direction over the switchcore interface (SCI);
- FIG. 9 is a simplified block diagram illustrating an embodiment of the controlled access ATM switch of the present invention in which there are an equal number of logic buffers in the input switchport and outlets for target switchports from the switchcore matrix;
- 25 FIG. 10 is a simplified block diagram of an embodiment of the controlled access ATM switch of the present invention in which a single logic buffer is used for all switchcore outlets to target switchports;
- FIG. 11 is a simplified block diagram illustrating buffer prioritizing and the use of variable buffer sizes in the input side of the switchports of the controlled access ATM switch of the present invention;
- FIG. 12 is a functional diagram showing the connection of the switchports to the switchcore and illustrating the principle employed for the access mechanism to the switch matrix in the controlled access ATM switch of the present invention;
- 30 FIG. 13 is a functional diagram illustrating a minimal solution for the positioning of buffers in the switchcore matrix which still provides the desired functionality of the controlled access ATM switch of the present invention;
- FIG. 14 is a functional diagram illustrating the positioning of buffers in the switchcore matrix when an intermediate number of buffers are employed in the controlled access ATM switch of the present invention;
- FIG. 15 is a functional diagram illustrating the positioning of buffers in the switchcore matrix in a complete solution in which one buffer, one cell deep, is used for each matrix cross-point in the controlled access ATM switch of the present invention;
- 35 FIG. 16 is a top level block diagram of a switchcore matrix;
- FIG. 17 is a block diagram of a row function unit (RFU) of the switchcore matrix;
- FIG. 18 is a block diagram of a column function unit (CFU) within the switchcore matrix;
- FIG. 19 is a high level flow chart of a software program which controls the functions of a CFU controller within the column function unit (CFU);
- 40 FIG. 20 is a block diagram of a cross function unit (XFU) within the switchcore matrix;
- FIG. 21 illustrates the timing relationship between the switchport-to-switchcore cell stream and the switchcore-to-switchport cell stream within the switchcore interface (SCI) for a specific switchport;
- FIG. 22 is a perspective view of one embodiment of the controlled access ATM switch of the present invention in which a single-chip switchcore is mounted on a back plane to which switchport boards are connected; and
- 45 FIG. 23 is a logical block diagram of the devices used for speed conversion in the preferred embodiment of the present invention.

## DETAILED DESCRIPTION OF THE INVENTION

[0017] FIG. 1 is a simplified block diagram of the major components of an controlled access ATM switch 10 of the type employed in the implementation of the principles of the present invention. The ATM switch of the present invention is essentially comprised of two parts: one or more switchports 11 and a switchcore 12. Each switchport 11 performs the logic switching and buffering functions of the ATM switch 10, and is divided into an input side 11A and an output side 11B, as shown in FIG. 2. The switchcore 12 essentially performs only routing functions based on a label added to each information cell, and may be implemented in a single chip. This greatly reduces the hardware and maintenance cost of the switchcore, and greatly increases its reliability, especially if a second plane 13 is provided for redundancy, as shown in FIG. 1.

[0018] FIG. 2 is a simplified functional block diagram illustrating the data flow between switchports 11 and the switchcore 12, and showing where input buffers 14 are placed in the switchports 11. In one embodiment, the switchcore 12 in the present invention is greatly simplified over the prior art switches. This is in part because the input buffers 14 are

5 moved from within the switchcore 12, where they are located in most conventional ATM switches, to the input side of the switchports 11A. The switchcore 12 therefore comprises only a switch matrix and a status register for each column in the matrix with an appended prioritizing mask register for fair transmission of cells. The switchcore 12 has a buffer depth of only one store for each path or route, and has no processor or any other communication channel for operation and maintenance.

10 [0019] The consequence of buffering in the switchport 11 is that the bandwidth must be increased on the input side 11A of each individual switchport connection to enable the switchport to receive a large amount of data in a short period or to receive data from several input sources nearly simultaneously. However, the storage capacity of each buffer 14 can be tailored to the individual requirements of each switchport 11, and of the type of traffic it handles, thereby improving overall system characteristics and cost. Each switchport 11 may be designed with a different sized input buffer 14, and may be made either more or less complex, depending on the type of communications service involved. For example, for conventional telephone service, with a continuous, low data-rate flow of information through the switch, a small input buffer 14 will suffice. For data communications, the information flow of which is more burst-like in nature, a larger input buffer 14 is required. Thus, overall cost of the ATM switch 10 is reduced since the design of the input buffers 14 can be tailored to the anticipated service type, and only the necessary buffer space provided.

15 [0020] In addition to the advantages of tailored individual buffer distribution, the introduction of access control and buffering in the switchport 11 provides a "guarantee" that the cells will pass through the switchcore 12 because the switchport 11 holds the cell in its buffer 14 while determining the status of its assigned routes through the switchcore 12, and releases the cell only when a route and a destination switchport is free.

20 [0021] The dotted line 15 in FIG. 2 represents the flow of access monitoring and protocol control information (PCI) which determines when each switchport 11 can transmit the information in its input buffer 14 through the switchcore 12. This monitoring and control process will be discussed in greater detail below.

25 [0022] The communication between the switchcore 12 and the different switchports 11 may be asynchronous or synchronous. The asynchronous communication is controlled by each switchport 11 and allows one switchport to send and receive at a high rate while another switchport sends and receives at a low rate. The synchronous communication requires that the switchports 11 requiring synchronization use a clock distribution signal in the switchcore 12. In this case one switchport acts as a master, and the other switchports act as slaves. The master switchport delivers the synchronizing clock signal to the slave switchports.

30 [0023] FIG. 3 is a simplified functional block diagram illustrating the manner in which operation and maintenance functions are monitored and controlled from the switchports 11. Each switchport 11 controls and monitors the operation and maintenance (O&M) functions on the routes in the switchcore 12 which are capable of being used by each respective switchport 11. The dashed lines 16a in FIG. 3 indicate that the O&M functions are performed on the routes through the switchcore 12, but not on the switchcore itself.

35 [0024] FIG. 4 is a simplified block diagram illustrating how all communication devices 17 connected to the controlled access ATM switch 10 have access to the switchcore 12 via the switchports 11. The switchports 11 provide the interface between the communication devices 17 and the switchcore 12. The switchports 11 may, for example, when carrying standard telephone traffic, convert the signal from standard C1 carrier format into ATM packet format having up to 56 bytes of information in each cell.

40 [0025] FIG. 5 is a simplified functional block diagram illustrating the logic switching and space switching functions performed by the controlled access ATM switch of the present invention. The logic and space switching functions are implemented through three protocol levels: ATM Logic Switching (ALS) 18, ATM Space Switching (ASS) 19 and physical framing 21. ALS 18 provides the interface between the ATM switch 10 and external network devices 17, and is performed within the switchports 11. ALS 18 translates incoming Virtual Channel Identifier/Virtual Path Identifier (VCI/VPI) numbers to outgoing numbers on both the input and output sides of the ATM switch 10. ASS 19 is a protocol which passes information cells between switchports 11 and the switchcore 12. ASS 19 is performed as the logic address from the switchport 11 is translated in the switchcore 12 to a physical address for space switching. Physical framing 21 indicates that the transfer of cells may be aligned with a framing reference in order to synchronize switchports 11 to each other.

45 [0026] Of the three protocols described, ATM Space Switching (ASS) 19 is the protocol with the greatest significance to the present invention. ASS 19 is a collection of functions and procedures carried out on an ATM Space Switch level. The functions are carried out partly in the switchport 11 and partly in the switchcore 12. The functions enable the extraction of cellsync and bytesync, maintenance of the switchcore 12, control of access to the switchcore, and determination of the status of sent/received cells. The functions are driven by the protocol control information (PCI) 15 transmitted from switchport 11 to switchcore 12 (and vice versa) and by the primitives from superior or controlling layers.

50 [0027] FIG. 6 is a simplified functional block diagram illustrating the logic and space switching functions when several controlled access ATM switches 10 are link-coupled in a matrix architecture or any other structure such as CLOS. Providing large buffers 14 in the switchports (SWP) 11a-d permits a high degree of concentration without impairing the properties of the system for data communication traffic through the switchcores (SWC) 12a-c. The divided line in the

ALS level of the inner switchports 11b and 11c illustrates symbolically that there are two mutually facing selector ports.

[0028] FIG. 7 is a simplified functional block diagram illustrating the structure and relationships between the three communications protocol levels which perform the logic and space switching functions in the controlled access ATM switch 10 of the present invention. An ATM-cell 24, comprising 53 bytes of information, may be stored in the ALS-PDU.

5 When the cell is put into the ALS-SDU 22, which comprises 56 bytes of information, three (3) bytes are left for free use. The ALS-SDU 22, together with the PCI 15 of 4 bytes, are then put into the ASS-PDU 27 which comprises 60 bytes, and are then relayed to the other ALS-entity 18 or vice versa. The functions are performed partly in the switchport 11 and partly in the switchcore 12. The functions are driven by the PCI 15 transmitted from switchport 11 to switchcore 12 (and vice versa) and by primitives from superior of controlling layers.

10 [0029] The switchcore interface (SCI) is the interface between the switchports 11 and the switchcore 12. Information cells, maintenance cells and idle cells are mixed on the SCI. The information cells are routed through the switchcore 12 while the maintenance and idle cells are terminated on both sides of the SCI.

15 [0030] FIG. 8 is a block diagram illustrating the physical lines of the SCI between one switchport 11 and one plane of the switchcore 12. The physical lines comprise a bidirectional CLOCK ref line 28, a DCLOCK SWP-SWC line 29 from switchport 11 to switchcore 12, a DATA SWP-SWC line 30 from switchport 11 to switchcore 12, a DATA SWC-SWP line 31 from switchcore 12 to switchport 11, and a DCLOCK SWC-SWP line 32 from switchcore 12 to switchport 11. Thus, each line except the CLOCK ref line 28 is implemented as a balanced pair.

20 [0031] FIG. 8a is a byte map of a generic cell 101 as it is sent in each direction over the SCI. The cell contains 60 bytes with bit 8 on byte 1 transmitted first in a serial bit-stream. Bytes 1-4 constitute an address and validation field 102, and bytes 5-60 are the payload (information) 103 carried by the cell 101. As an option for high data rates, and in particular when optical transmission line is used, a Line Code Bit (LCB) 104 may be inserted every 24th bit. Together with a two-step scrambling, the LCB 104 gives good DC balance. The switchcore 12 detects the LCB 104 and uses the same technique in the opposite direction for each individual switchport 11.

25 [0032] A cell type field (CTF) 105 is a two-bit coded field which is used in both directions. The codes in the CTF 105 indicate what type of cell is being transferred. The following codes, with their interpreted meanings are included:

Code	Type of Cell	Remarks
00	Idle cell	RAF not valid; RPF valid.
01	Maintenance	Carries maintenance command; RAF, RPF replaced by maintenance fields; see maintenance cell format (FIG. 8c).
10	Active traffic	Low prio cell; RAF, RPF valid
11	Active traffic	High prio; RAF, RPF valid.

30 [0033] A tag error check (TEC) field 106, comprising a field of 6 bits, is generated and checked on both sides of the SCI. The TEC 106 is used for both cell synchronization and validation of the previous 26 bits in the cell.

40 [0034] FIG. 8b is a byte map of an information (traffic) cell 111 as it is sent in each direction over the SCI. Bytes 1-3 of the information cell 111 comprise a bitmap pinpointing the individual switchports 11 on the outside of the switchcore 12. In the sending direction (switchport to switchcore), bytes 1-3 comprise the relay address field (RAF) 25 where each bit indicates a target (receiving) switchport on the other side of the switchcore 12. In the receiving direction (switchcore to switchport), bytes 1-3 comprise the relay poll field (RPF) 26, and indicate which target switchports are occupied and which are free.

45 [0035] FIG. 8c is a byte map of a maintenance cell 121 as it is sent in each direction over the SCI. A number of maintenance commands may be issued from a switchport 11 to the switchcore 12 concerning the parts of the switchcore 12 which correspond to the sending switchport 11. In the sending direction, byte 1 contains a two-bit rate data field (RDF) 122 in which the following codes are included:

Code	Meaning
00	Any rate difference between own switchport and addressed switchport can be accommodated.
01	Own switchport sending rate is higher than receiving rate from the addressed switchport.
10	Own switchport receiving rate is higher than sending rate from the addressed switchport.

**EP 1 051 001 A2**

(continued)

Code	Meaning
11	Own switchport rate is synchronized with addressed switchport.

**[0036]** In the sending direction, byte 2 contains a six-bit switchport address field (SPAF) 123 which provides the address of the switchport 11 which is sending the maintenance cell 121. There are 24 switchports 11, numbered 0-23 binary.

**[0037]** In the sending direction, byte 3 contains a two-bit plane select field (PLSF) 124 which selects which of the switchcore planes is to carry out the maintenance command. The following commands are included:

Code	Meaning
00	The command is not carried out.
01	Only plane A carries out the command; both planes send acknowledgment.
10	Only plane B carries out the command; both planes send acknowledgment.
11	Planes A and B carry out the command.

**[0038]** Byte 3 also contains a four-bit operation request field (ORF) 125. The ORF 125 may be used to request such actions as block or unblock an addressed switchport, open or close an addressed clock reference gate, set rate data between own and addressed switchport, set throttling for own switchport, clear own column or row, and set switchcore internal cell delay. The switchcore internal delay command may be used to set the delay in the switchcore so an isochronal serial transfer of cells can be made between attached devices without unnecessary delay losses. Serial isochronal support requires minimum delay is the attached devices while variable cell traffic support requires maximum delay in the switchport in order to analyze the buffer situation.

**[0039]** In the receiving direction, byte 3 contains a two-bit operation indication field (OIF) 126 which indicates to the switchport 11 the status of the previous cell from the switchport 11 to the switchcore 12. The OIF 126 indicates whether the previous cell had an error, or in case of a maintenance cell to the switchcore 12, whether or not it was carried out. The following codes are included:

Code	Meaning
00	Not used.
01	Previous switchport to switchcore maintenance command carried out.
10	TEC-error in previous cell.
11	Error in field interpretation of previous switchport to switchcore maintenance cell.

**[0040]** In the receiving direction, byte 5 contains a five-bit switchport identification number 127 indicating what number the switchcore 12 has given the switchport 11. The switchport identification number 127 corresponds to the SPAF 123 in the sending direction.

**[0041]** Byte 5 also contains a one-bit synchronization window field (W) 128 which indicates the size of the synchronization window. The synchronization window is described in greater detail below. The following codes are included:

Code	Meaning
00 (default)	Window corresponding to the timing of byte 2, allowing an additional time corresponding to one byte for CLOCK ref and switchcore internal jitter.
01	Window is 60 bytes.

[0042] The switchcore also sends its own article number and revision number in bytes 6 through 9, as shown by shaded area 129 in FIG. 8c.

[0043] FIG. 8d is a byte map of an idle cell 141 as it is sent in each direction over the SCI. The idle cell 141 is identical to the maintenance cell 121 except that the first three bytes 142 of the idle cell 141 in the sending direction (switchport to switchcore) has no significance.

[0044] FIG. 9 is a simplified block diagram illustrating key elements of the preferred embodiment of the controlled access ATM switch of the present invention in which there are an equal number of logic buffers 14 in the input switchport 11a, and outlets for target switchports 11b from the switchcore matrix 12. In the illustrated case, there are "n" input switchports 11a and "n" output switchports 11b. In the preferred embodiment, the switchcore comprises a switch matrix 12 with n rows 12a and n columns 12b. Accordingly, there are n inlet switchports 11a and n outlet switchports 11b.

[0045] FIG. 9 has been simplified to illustrate only a single inlet switchport 11a (SWP No. X), a single row 12a within the switchcore 12, and n outlet switchports 11b (SWP 1-n). Each inlet switchport 11a is connected to a different row of the switchcore 12. For example, inlet switchport SWP No. 1 (not shown) is connected to row no. 1 of the switchcore matrix 12, inlet switchport SWP No. 2 is connected to row no. 2, inlet switchport SWP No. 3 to row no. 3, etc. In FIG. 9, inlet switchport SWP No. X is connected to row X (12a). Each of the outlet switchports SWP 1-n are shown to be connected to the same matrix row 12a in the switchcore 12. In this particular case the outlet switchports are connected to row X to which inlet switchport SWP No. X is connected. Each other inlet switchport, not shown in FIG. 9, is connected to its respective matrix row, not shown in FIG. 9, and each of the outlet switchports 11b shown in FIG. 9 has additional connections, not shown in FIG. 9, to each of the other rows of the switchcore matrix 12.

[0046] Each inlet switchport 11a includes n input logic buffers 14, all connected to a single row of the switchcore matrix 12. The row differs for each inlet switchport 11a. The single inlet switchport illustrated in FIG. 9 (SWP No. X) is shown to include n input buffers 14, all connected to row X. The input logic buffers 14 store and retrieve information cells on a first-in, first-out (FIFO) basis, and are labeled "FIFO 1" through "FIFO n" in FIG. 9. Each buffer 14 is physically mapped on its corresponding switchport 11.

[0047] In the case of register one (r1), mapping is effected from a logic buffer number to a physical buffer number. This means, for example, that buffer n-1 (FIFO n-1) will land in r1, position n-1, and FIFO n will land in r1, position n. A cell arriving at an inlet switchport 11a has its destination address field examined and is stored in the FIFO buffer 14 corresponding to the physical destination of the cell address, i.e., the cell is temporarily stored in a FIFO or queue that goes to the destination switchport 11b. In FIG. 9 legends r1, r2, and r3 denote registers provided at the inlet switchport 11a. Each register r1-r3 has a number of bit positions corresponding to the total number of buffers 14, with each position corresponding to a particular buffer. These registers r1-r3 are used to physically map each buffer 14 on its corresponding switchport 11a. Each other inlet switchport 11a not shown in FIG. 9 has corresponding registers.

[0048] Register r1 may be utilized, for example, by placing a one (1) in those bit positions corresponding to buffers 14 that contain information cells. Bit positions in r1 for which the corresponding buffer is empty may be indicated by a zero (0). In FIG. 9, bit positions of register r1 that are shown as shaded indicate that an information cell is currently in the corresponding FIFO buffer 14, and bit position corresponding to empty buffers are shown in white.

[0049] Register r2 contains the latest status of the receivers (target switchports) 11b on the other side of the switchcore 12, i.e., the content of RPF 26. Each bit position represents a target switchport 11b. Specifically, bit position 1 of register r2 corresponds to the outlet switchport 11b to which the cell in bit position 1 of register r1 is addressed; bit position 2 of register r2 corresponds to the outlet switchport to which the cell in bit position 2 of register r2 is addressed, etc. A shaded bit position indicates that the target switchport 11b is free, and a white bit position indicates the target switchport is occupied. By bit-wise AND-ing the contents of register r1 and the contents of register r2 (i.e., r3 = r1 AND r2), and storing the result in register r3, register r3 indicates, in (shaded) register positions that store a one (1), only those FIFO buffers 14 containing cells which can be received at switchports 11b that are ready to receive. In the example given in FIG. 9, the contents of register r1 indicates that FIFO 1, FIFO 3, FIFO 4 and FIFO n-1 have cells to send. The contents of register r2 indicates that outlet switchports SWP 1, SWP 4 and SWP n are ready to receive cells. The result of the AND-operation, i.e. the contents of register r3, indicates that only FIFO 1 and FIFO 4 are allowed to send their respective cells. FIFO 3 is denied to send its cell, because the status of its target switchport 11b, represented by the zero (0) at bit position 3 of register r2, indicates that the target switchport 11b is not ready to receive any new cells.

[0050] In order to utilize the switching capacity of the access controlled ATM switching system of the present invention to 100 percent and, at the same time, ensure that a buffer cannot be totally excluded, a rotary priority indicator (labeled "prio-pointer") is implemented. Cells from two different FIFO buffers 14 cannot be sent through the switchcore simultaneously because both cells are transferred on the same switchcore matrix row 12a. The priority indicator, therefore, gives priority to only one FIFO. In FIG. 9, the priority indicator points at FIFO 3 in register r1. Software controlling the priority pointer, however, may give priority to FIFO 4 since the target switchport 11b, as indicated by register r3, is occupied. FIFO 4 sends its cell, as is schematically indicated by arrow 23, to the indicated matrix row 12a. The address field of the cell will again be read, and the addressed outlet switchport 11b extracts the cell and transfers it to its associated device 17 (FIG. 4).

[0051] The outlet switchport 11b which extracts the cell at first sets a flag in the RPF 26 of the extracted cell and then sends the extracted cell to its target switchport. The flag set in the RPF 26 indicates that the outlet switchport 11b is occupied with receiving a cell. In this particular case outlet switchport SWP 4 extracts the cell, sets the flag, and forwards the cell to its associated device 17. Finally, FIFO 1 is allowed to send its cell.

[0052] The aforesaid is only one illustrative embodiment of a method for structuring buffers and analyzing which cell to transmit next, and other methods may be implemented. The illustrated method can be performed within a period of one microsecond with some simple operation in, for example, a risc-processor. With the aid of specific hardware, an analysis speed of less than 200 nanoseconds is possible.

[0053] FIG. 10 is a simplified block diagram of an embodiment of the controlled access ATM switch of the present invention in which a single logic buffer 14 is used for all switchcore outlets to target switchports 11b. In many applications such as Switched Multimegabit Data Service (SMDS), a single input buffer 14 will suffice, irrespective of the addressed output on the other side of the switchcore 12. In SMDS, the main traffic always passes from one switchport 11 to another for the capacity-critical paths when concentrating from several accesses to a server.

[0054] The single buffer 14 may address a single target switchport 11b, or it may group-address several target switchports 11b. A simple two-stage process for group-addressing is shown in FIG. 10. In stage 1, register r1 indicates in black the target switchports 11b to which a cell in the buffer 14 is to be sent, in this example, switchports 1, 3, 4, and n-1. Register r2 indicates in black the target switchports 11b that are free to receive the next cell (RPF 26), in this example, switchports 1, 4, and n. Register r3 results from the operation "r3 = r1 AND r2" and thus indicates the target switchports 11b to be addressed (RAF 25) in stage 1, in this example, switchports 1 and 4.

[0055] In stage two, all remaining group addressees (switchports 3 and n-1) are expedited, as shown in black in register r1. Register r2 again illustrates which target switchports 11b are free to receive (2, 3, 4, and n-1). After the operation "r3 = r1 AND r2", register r3 shows that cells are to be sent to target switchports 3 and n-1. If target switchports 3 and/or n-1 are not free, the procedure is repeated until the cell has been passed to all of the group addressees.

[0056] FIG. 11 is a simplified block diagram illustrating buffer prioritizing and the use of variable buffer sizes in the input side of the switchports 11A of the controlled access ATM switch of the present invention. Buffers 14 of differing capacities can be utilized, depending on the type of communications traffic concerned. FIG. 11 illustrates the differing buffer sizes between a buffer for Variable Bit Rate (VBR) traffic 35 and a buffer for Constant Bit Rate (CBR) traffic 36, where CBR traffic has been assumed to require less buffer capacity.

[0057] FIG. 11 also illustrates that a method for prioritizing the information from each buffer may also be implemented in the ATM switch 10. The high priority (HPRIOR) block 37 represents a method to, for example, provide higher priority for the information from the CBR buffer 36. The buffering and prioritizing functions are fully implemented in the input side of the switchports 11A, and are optimized for the type of communications service concerned.

[0058] FIG. 12 is a functional diagram illustrating the connection of the switchports 11 to the switchcore 12 and the principle employed within the access mechanism to the switch matrix in the controlled access ATM switch of the present invention. The switchcore 12 is comprised of a switch matrix represented in FIG. 12 as rows R1 through Rn and columns C1 through Cn. The rows represent inputs from input switchports 11a, and the columns represent outputs to target switchports 11b. At the points in the switch matrix where the row number and column number are equal, the input side of the corresponding switchport 11A will transmit a cell to its own output side 11B. For example, at the intersection of row 1 and column 1, the input side 11A of switchport (SWP) 1 transmits cells to row 1, and column 1 then transmits the cell in column 1 to the output side 11B of switchport 1.

[0059] The connection of the switchports 11 to the switchcore 12 and the principle employed within the access mechanism is based on phase shifting of incoming and outgoing cells. The extent of the phase shift depends on the length of time taken to process and assemble RAF 25 and RPF 26 using the method illustrated in FIGS. 9 and 10. Possible series/parallel conversions may also take time.

[0060] FIG. 12 also illustrates how RAF 25 and RPF 26 can appear to the first switchport (SWP 1) at different times. At time  $t_0$ , SWP 1 receives RPF 26 which identifies all target switchports 11b which are free to receive cells. The switchport then compares the RPF 26 with the incoming RAF 25 which identifies which target switchports 11b are addressed by the cell in the SWP 1 buffer 14. This comparison consists of a simple AND function, and is illustrated in FIG. 12 by dotted lines 41 and 42. This comparison is completed at time  $t_1$ , and identifies target switchports 1 and 4. SWP 1 transmits the addressed cell to row R1, and to columns 1 and 4 which correspond to those target switchports 11b to which the cell can be sent, in this example, SWPs 1 and 4. This transmission is illustrated in FIG. 12 by the dotted lines 43 and 44 leading from the RAF 25 to positions R1,C1 and R1,C4 in the switch matrix.

[0061] The next RPF 26 arrives at time  $t_2$ , and indicates that target switchport SWP 2 is free. SWP 1 then compares the RPF 26 to the incoming RAF 25 which indicates that SWP 2 is addressed by the cell in the SWP 1 buffer 14. This comparison is illustrated by dotted line 45, and is completed at time  $t_3$  when the cell is transmitted to position R1,C2, the position corresponding to SWP 2. This transmission is illustrated by dotted line 46 leading from the RAF 25 to position R1,C2 in the switch matrix.

[0062] At time  $t_4$ , RPF 26 indicates that all target switchports 11b are free to receive cells. However, at time  $t_5$ , the

incoming RAF 25 indicates that SWP 1 has no addressed cells to send, and therefore, the AND comparison results in no cells being transmitted.

[0063] As noted above, the switchcore matrix 12 has a buffer depth of only one cell for each path or route. The buffers in the switchcore 12 may be implemented in one of several ways, ranging from a minimal solution to a complete solution with a buffer at each crosspoint of the matrix.

[0064] FIG. 13 is a functional diagram illustrating a minimal solution for the positioning of buffers 51 in the switchcore matrix 12 which still provides the desired functionality of the controlled access ATM switch of the present invention. Even this minimal solution, however, provides sufficient switch performance for services such as Switched Multimegabit Data Service (SMDS).

[0065] The minimal solution of FIG. 13 provides a "pool" of common buffers at the input of the switchcore 12. Provided that a buffer 51 is free, the switchcore 12 will signal the corresponding switchport 11 that the switchcore 12 is able to receive a new cell. Each block 51 in FIG. 13 represents from one to twelve buffers organized as a shared pool of buffers. The number of buffers 51 may vary, but twelve is the useful maximum because the peripheral logic grows to such an extent that, for more than twelve buffers 51, it becomes more economical to spread the buffers on each crosspoint of the matrix. The common buffer pool may also be distributed across the switchcore matrix 12 to those crosspoints which are used most often.

[0066] FIG. 14 is a functional diagram illustrating the position of buffers 51 in the switchcore matrix 12 when an intermediate number of buffers are employed in the controlled access ATM switch of the present invention. FIG. 14 illustrates a solution in which each buffer 51 is shared by two crosspoints in the matrix 12, but other divisions are also possible within the scope of the present invention.

[0067] FIG. 15 is a functional diagram illustrating the position of buffers 51 in the switchcore matrix 12 in a complete solution in which one buffer, one cell deep, is used for each matrix cross-point in the controlled access ATM switch of the present invention. Other solutions are possible, depending on chip layout and other physical limitation, and remain within the scope of the present invention. In one embodiment, a 20x20 matrix 12 and one buffer 51 for each crosspoint results in an approximate memory capacity of 179,200 bits, divided on 400 buffers of 56x8.

[0068] FIG. 16 is a top level block diagram of a switchcore matrix 12. The switchcore 12 comprises three basic units for each switchport, i.e., 24 of each unit in the preferred embodiment. On a per-switchport basis, a row function unit (RFU) 61 terminates the incoming cell stream 62. A column function unit (CFU) 63 forms a synchronized pair with the RFU 61, and generates the outgoing cell stream 64. A cross function unit (XFU) 65 receives information cells 111 (FIG. 8a) from the RFU 61 via the row bus 66 and relays the information cells through the switchcore 12. The RFU 61 throws away idle cells 141 (FIG. 8d), and decodes and executes maintenance cells 121 (FIG. 8c).

[0069] Each CFU 63 searches the XFUs 65 attached to the CFU for cells to be relayed, and extracts those cells via a column bus 67. If no cells are found, the CFU 63 generates an idle cell 141 which is transmitted to the attached switchport 11. If an incoming maintenance cell 121 is detected, then the stated command is executed and an acknowledgement is sent to the switchport 11. If any field is out of range, an error acknowledgement will be sent instead.

[0070] Each XFU 65 stores an addressed cell, and sets a flag indicating that a cell is waiting to be unloaded by the CFU 63.

[0071] FIG. 17 is a block diagram of a row function unit (RFU) 61 of the switchcore matrix 12. It can be seen that the RFU 61 interfaces with the switchport 11, the column bus 67 and row bus 66, and the CFU 63. A phase aligner 71 adapts to the incoming bit rate that may vary from a very low speed (a few bit/s) up to the technology limit which may be approximately 200 Mbit/s, and aligns the incoming bit rate with the incoming clock. A cell framer 72 performs the function of converting the incoming bitstream into byte format and finding the start of a cell in order to synchronize the other internal units in the RFU 61 as well as the associated CFU 63 and all XFUs 65 attached to the RFU-CFU pair. The RFU 61 uses the tag error check (TEC) 106 in order to find the start of the cell. A line code ejector 73 may comprise a 5-bit modulo 25 counter that removes a line code polarity bit from the data stream by prolonging every third byte with the time of the line code bit. A RFU controller 74 derives the plane select field (PLSF) 124 (FIG. 8c), the operation request field (ORF) 125, and the cell type field (CTF) 105 and stores their values at the times they are present on the cell data bus. At designated times, the PLSF 124, ORF 125, and CTF 105 are sent over the row bus 66 to the CFU 63. The clock buffer 75 is a bidirectional buffer controlled by the RFU controller 74.

[0072] FIG. 18 is a block diagram of a column function unit (CFU) 63 within the switchcore matrix 12. The CFU 63 interfaces with the column bus 67 (FIG. 16), with the RFU 61, and with the switchport 11. When the CFU 63 receives a cell-sync signal from the RFU 61, indicating that a cell addressed to that CFU has been sent to a cross function unit (XFU) 65, the CFU 63 unloads the cell from the XFU 65 via the column bus 67. If there is no cell, the CFU 63 generates an idle cell 141 (FIG. 8d). If the RFU 61 indicates that a maintenance cell 121 (FIG. 8c) has been sent, the CFU 63 generates a maintenance cell 121. The unloaded cell, the idle cell, or the maintenance cell is added to the relay poll field (RPF) 26 and sent to the switchport 11 along with a clock signal indicating a valid bit.

[0073] A CFU controller 81 controls the actions of the CFU 63. FIG. 19 is a high level flow chart of the software program which controls the functions of the CFU controller 81. The program is entered at step 82 when the flow of

cells begins. At step 83, the CFU 63 receives a cellsync indication from the RFU 61 indicating that a cell has been received which is addressed to the CFU 63. At step 84, it is determined whether or not a maintenance command is present. If a maintenance command is present, the program moves to step 85 where the CFU 63 carries out the maintenance command. At step 86, the CFU 63 generates a maintenance cell 121 (FIG. 8c). If, however, at step 84 it was determined that a maintenance command was not present, then the program moves to step 87 where a scan is performed in an attempt to retrieve a cell from the XFU 65. At step 88, it is determined whether or not a cell was found on the scan. If a cell was not found, the program moves to step 89 where an idle cell 141 is generated. If, however, at step 88 a cell was found, then the program moves to step 90 where the cell is unloaded from the XFU 65.

**[0074]** Referring again to FIG. 18, an idle cell generator 91 generates bits 5 to 60 of an outgoing idle and maintenance cell upon command from the CFU controller 81. A cell assembly device 92 assembles cells in the formats shown in FIGS. 8a-8d. The first three bytes are generally poll data, and the fourth byte contains cell type field (CTF) 105 and tag error check (TEC) 106. All data in the first four bytes except the TEC 106 are only put in the byte stream by control signals from the CFU controller 81. In addition, the payload 103 is loaded, which may be an idle or maintenance pattern or an information cell unloaded from an XFU 65. A delay line of 8 bytes is inserted in front of the payload 103 for a late arriving poll result.

**[0075]** A priority device 93 supports the CFU controller 81 by storing the results of scans when the CFU 63 scans XFUs 65 for loaded cells. The priority device 93 indicates a found cell and provides the CFU controller 81 with the selected XFU address. If the scan result is negative, i.e., there were no cells to relay, a miss is indicated to the CFU controller 81.

**[0076]** A throttle device 94 is shown in phantom in FIG. 18, and is an optional device. The throttle device may be, for example, a settleable modulo counter of 5 bits. It allows a connected switchport 11 to get a much lower logical throughput than the actual physical rate allows.

**[0077]** A CFU line code ejector 95 inserts a line code bit every 25th bit when so indicated by the RFU 61. A CFU phase aligner 96 adapts to the same clock and levels as the RFU phase aligner 71 (FIG. 17). In addition, the CFU phase aligner 96 converts parallel data to a serial bit stream.

**[0078]** FIG. 20 is a block diagram of a cross function unit (XFU) 65 within the switchcore matrix 12. The XFU 65 interfaces with the row bus 66 and the column bus 67 (FIG. 16). The cells on the row bus 66 are written into an XFU memory device 151 when the relay address field (RAF) 25 matches the XFU row address. The cells are unloaded from the XFU 65 onto the column bus 67 if the XFU in question is addressed. In addition, current XFU status is read by polling the row bus 66. Current XFU status can be read by scanning from the CFU 63.

**[0079]** The XFU 65 is controlled by an XFU controller 152 which decodes incoming signals from the RFU 61 on the row bus 66, and incoming signals from the CFU 63 on the column bus 67. An input logic device 153 analyzes the 24-bit relay address field (RAF) 25 in incoming cells on the row bus 66 to determine if the XFU in question is addressed. An output logic device 154 determines when the XFU 65 is addressed by control lines on the column bus 67. A clock gate device 155 consists of one flip-flop and a gate that is enabled by the flip-flop. The flip-flop is in a reset state, and the clock gate is disabled, at power up.

**[0080]** The XFU memory device 151 may be a two-port memory implemented as a register file with three state outputs. This implementation is consistent with a gate-array embodiment of the switchcore matrix 12.

**[0081]** FIG. 21 illustrates the timing relationship between the switchport-to-switchcore cell stream and the switchcore-to-switchport cell stream within the switchcore interface (SCI) for a specific switchport 11. In FIG. 21, and the preferred embodiment, the two streams of cells are synchronized at byte 20 of the SWP-to-SWC stream 62. The processing time in the switchcore 12 sets the actual synchronization time and varies for different switchcore embodiments. The SWC-to-SWP stream 64 follows the SWP-to-SWC stream 62 by a time period sufficient for the switchport 11 to analyze the RPF field 26 and determine if the next cell can be addressed to the targeted switchport, and thereby construct the RAF 25 for the next cell. The time period for the switchport to perform this analysis and send the next cell is indicated by the "association" arrow in FIG. 21, and in the preferred embodiment, is the time it takes to transmit 40 bytes. The "delay" arrow indicates the delay which may be set by the delay command in the operation request field (ORF) 125 (FIG. 8c). The "previous" arrow indicates the relationship between the operation indication field (OIF) 126 (FIG. 8c) and the RAF 25 of the previous SWP-to-SWC cell. The OIF 126 indicates whether the previous cell had an error, or in case of a maintenance cell to the switchcore 12, whether or not it was carried out.

**[0082]** A clock reference signal is generated through the SCI in order to achieve cell synchronization (cell-sync). In the preferred embodiment, the switchports 11 are synchronized so that their cell starts fall within a window corresponding to the time it takes to transmit two bytes in order to utilize the full throughput. An additional time period corresponding to the time to transmit approximately one bit is added to the window to account for clock reference jitter. The system also allows for an additional half-byte time period for internal jitter in the switchcore 12.

**[0083]** Redundancy may be added to the ATM switch 10 of the present invention in several ways. For example, a second plane 13 may be added to the switchcore 12 as shown in FIG. 1. The planes 12 and 13 may be mutually asynchronous, depending on the difficulty of synchronizing the switch with possible loss of a cell in the switchcore 12.

Adding an asynchronous plane 13 adds to the expense and complexity of the switchport 11 because the switchport 11 must be made more intelligent with several measuring algorithms.

[0084] The switchcore 12 of the controlled access ATM switch 10 of the present invention may be constructed on a single chip which has the capacity for 20 double-directed 155-Mbit connections, the buffers 51, and remaining switchcore functions. Such a single-chip switchcore 12 may be mounted directly on a back plane which is not much wider than the ATM switch 10 itself. FIG. 22 is a perspective view of one embodiment of the controlled access ATM switch 10 of the present invention in which each single-chip switchcore plane 12 and 13 is mounted on respective back plane strips 161 and 162 to which switchport boards 11 are connected. The back plane strips 161 and 162 are replaceable, just as other circuit boards are replaceable. The back plane strips 161 and 162 need not be straight; they may also be curved or folded through 90° since only one switchcore is needed for connecting the switchport boards 11 on each strip.

[0085] If it is wished to maintain a lower level of technology, the switchcore 12 may be divided into three or four chips with a corresponding reduction in speed and required internal memory size. Conversely, four switchcore chips of 155-Mbit capacity may be link-coupled together in order to upgrade the switchcore to 620-Mbit capacity. Link-coupling requires the installation of switchports 11 between each switchcore 12. Therefore, for a number of link-coupled structures, the ATM switch 10 cannot be of plane-duplicated construction. From a reliability standpoint, this need not be a disadvantage. A plane-duplicated switch, in essence, is a switch with  $n+1$  redundancy, where  $n=1$ . There are a number of link-coupled structures which provide  $n+1$  redundancy, where  $n$  is greater than 1 in various stages of the structure.

[0086] Another advantage of the controlled access ATM switch 10 of the present invention is that the built-in access control supports the connection of different devices which operate at different physical speeds. The ATM switch of the present invention provides for total asynchronous communications at any speed. Although the switchcore 12 may become slightly more complex and expensive as a result of this capability, benefits are obtained on the device (switchport) side which outweigh the additional switchcore expense and provide an improved overall cost profile.

[0087] Isochronal service through the switchcore 12 may be created by subordinating the ATM cells in a frame. With isochronal service, the controlled access ATM switch 10 of the present invention can handle both STM and ATM traffic, and may be used in multimedia terminals intended for such services as PABX and public access nodes.

[0088] Isochronal service is based on the ATM-cell format, although it is coupled through the switch 10 at a specific predetermined time relative to a subordinate 125- $\mu$ s frame. Due to the clock distribution signal, a master device attached to a switchport 11 can send its 125- $\mu$ s frame sync to slaved devices attached to other switchports. The devices then schedule their cells within the time frame given by the master. Thus, it is essential that no other time slot on the same column have requested isochronal output at the time at which the cell/time slot shall be read. Therefore, the switchports 11 must coordinate the scheduling of isochronal cells by, for example, a controlling administrator which distributes isochronal time slots on a column basis so that no cell collisions occur.

[0089] The administrator, which may be centralized or distributed, may also distribute isochronal time slots so that a minimum delay occurs in the switchcore 12, since the time spent by the cell in a switchcore buffer 51 wastes buffer resources. The smallest switching level is thus a cell, meaning that 3.6 Mbit is the lowest conceivable allocation bandwidth in a single frame structure then expediting isochronal service at a 125- $\mu$ s frame level.

[0090] The switch 10 may also use a multiframe structure which obviates the need of "sacrificing" a full cell in those cases when the incoming bandwidth is longer than 125  $\mu$ s. The frame or multiframe structure may be based either on a synchronization pattern in the data flow, or on a frame clock which may run, for example, at 8 KHz, or a combination of both. An 8-Khz clock may result in some jitter problems, but it will provide a less expensive hardware solution since it can be provided through the clock distribution.

[0091] For switching to a lower level than one cell, the controlled access ATM switch 10 may be equipped with a device which can switch 155-Mbit currents on a 64-kbit level (a 4/0 device). With a 4/0 device, the internal structure of the cells is dissolved, and bytes are moved between cells and are then transmitted in different directions.

[0092] A proposed standard for circuit emulation in an ATM-environment will probably eliminate the need for switching on the 64-kbit level when conversion to ATM occurs only once. The standard specifies that the ATM cell is allocated to a connection and is partially or fully filled with 64-kbit samples, thus making it possible to more efficiently utilize the switching capacity.

[0093] It should also be noted that the controlled access ATM switch 10 of the present invention will function very well in single-plane solutions by including error-discovery mechanisms in the cell as it is transferred through the switchcore 12 from one switchport 11 to another. Three additional bytes may be added for this purpose. This process would be difficult to achieve in a pure circuit switch without incurring greater expense. This capability renders the controlled access ATM switch 10 suitable as an access switch for multimedia applications.

##### 55 Flow control

[0094] The bandwidth may vary from device to device amongst the devices 17 from which an inlet switchport 11a may receive cells. Therefore, an inlet switchport may receive cells at a speed which is higher, lower, or equal to the

speed at which a target switchport 11b can output its cells to its associated device 17. This puts restrictions on the flow of cells through the ATM switch so as not to corrupt the data in the payload of the received cells. If, for example the speed, referred to as the output speed, by which a target switchport can output its cells to its associated device 17 is lower than the speed, referred to as the input speed, at which cells, which are addressed to the said target switchport, arrive at an inlet switchport, then the inlet switchport must be prevented from receiving and temporarily storing the following cell until the previous cell has been received by the outlet switchport. If this rule is violated, there is a high probability that the next cell will overwrite parts of the cell being extracted. The cell being extracted will then contain corrupted data. On the other hand, if the input and output speeds are equal, or if the input speed is lower than the output speed, then the inlet switchport can start to receive a new cell as soon as the target switchport starts to receive the previous cell with no risk for overwriting the cell being extracted.

[0095] In accordance with the present invention the above is achieved by using specified rules for setting the RPF flag and a scan status flag, referred to as the SS-flag. The use of flag rules makes it possible to achieve "speed conversion" in the switch. Speed conversion means that the bit rate at which cells arrive at the switch is different from the bit rate at which cells are leaving the switch.

[0096] The speed conversion feature affords the switch of the present invention a unique feature: the capability to replace one or more, but not all, of the devices 17, designed to operate with a first specific bit rate, with new devices designed to operate with a second specific bit rate, different from the first.

[0097] If, for example, each of the devices 17 in FIG. 4 represent devices that each terminate a respective trunk designed to operate at a bit rate of 64 kbit/s, one or more of the trunk terminating devices may be replaced with new devices designed to terminate trunks that operate at a bit rate of 155 Mbit/s without the need to replace all of the devices 17. Accordingly, just individual trunks, and not all trunks of a network structure, may be upgraded and still use the same switch. This is different from existing ATM switches which require changing all devices 17 and upgrading all the trunks of the network in order to retain the same ATM switch.

[0098] The RPF field 26 of a cell from the switchcore to a switchport comprises the status of each of the outlet switchports 11b of the switch at the time the RPF field 26 is composed. In other words, the RPF 26 comprises a snapshot of the status of all of the outlet switchports 11b, whether a specific switchport is ready to receive a new cell or not. An individual bit of the RPF field 26 can be compared to a traffic light showing a red or green light for the traffic to the specific outlet switchport represented by the bit. If the bit is a red light, for example, indicated by setting the bit to logical "1", the buffer 51 (FIG. 15) at the matrix cross point which corresponds to the outlet switchport that the bit in the RPF field represents, is not ready to receive a new cell. If the bit is a green light, indicated by setting the same bit to logical "0", the buffer 51 is ready to receive a further cell addressed to the same outlet switchport 11b.

[0099] FIG. 23 is a logical block diagram of the devices used for speed conversion in the preferred embodiment of the present invention. FIG. 23 relates to FIG. 15 wherein there is one buffer 51 at each matrix crosspoint. FIG. 23 discloses a set of devices associated with a single matrix crosspoint, for example 2,3, the crosspoint between row number 2 and column number 3. Unless indicated otherwise, each of the other crosspoints of the matrix are associated with a similar set of devices. FIG. 23 further illustrates speed conversion when a cell arriving at inlet switchport SWP2 is routed to outlet switchport SWP10. The inlet bit rate at SWP2 is denoted "i" and the outlet bit rate at SWP10 is denoted "o". Four cases must be considered, namely  $i > o$ ,  $i < o$ ,  $i = o$  and  $i \neq o$ . The notation  $i \neq o$  means that the difference between  $i$  and  $o$  is unknown which, for example, is the case upon initialization when the switch does not know what devices 17 are connected. A poll status register (PSR) 160 is reading (polling) the status of FIFO buffer 51 in order to see if the buffer 51 is ready to receive a next cell addressed to SWP10. This polling is schematically indicated by the dotted line 161. A scan status register (SSR) 162 is scanning the status of the same FIFO buffer 51 in order to see if the buffer contains any cell to send. This scanning is schematically indicated by the dotted line 163. Both the PSR 160 and the SSR 162 comprise a flip-flop circuit having a set and a reset input and an output.

[0100] Resetting of the PSR 160 is conditional, illustrated by a movable arrow 164 which chooses between two reset conditions, R1 and R2. The choice between R1 and R2 is made by a poll rate register (PRR) 165. The PRR 165 contains information relating to the difference between the bit rate of the device connected to matrix row 2 (SWP2) and the devices 17 connected to matrix column 3, in this case the devices connected to SWP 1-n (FIG. 9). This information is static and is loaded into the PRR 165 when the switch is initialized or when a specific device 17 is replaced with a new device operating at a higher or lower bit rate than the replaced device.

[0101] The information in the PSR 160 is dynamic and is changed each time a cell arrives in, or is extracted from, the FIFO buffer 51. The output signal 166 of the PSR 160 is an RPF bit relating to the output switchport 11b which is addressed by the cell currently stored in the FIFO buffer 51, e.g., SWP 10.

[0102] For the SSR 162, setting is conditional. Movable arrow 167 chooses between two set conditions, S1 and S2. The choice between S1 and S2 is made by a scan rate register (SRR) 168. The SRR 168 contains information relating to the difference between the bit rate of the device connected to matrix row 2 (SWP 2) and the devices 17 connected to matrix column 3, in this case the devices connected to SWP 1-n (FIG. 9). Like the information in the PRR 165 this information is also static and is loaded into the SRR 168 when the switch is initialized or when a specific device 17 is

replaced with a new device operating at a higher or lower bit rate than the replaced device.

[0103] As noted above, the contents of the scan status register (SSR) 162 is dynamic and is changed each time a cell arrives in, or is extracted from, the FIFO buffer 51. The output 169 of the SSR 162 includes a scan status flag (SS-flag). The output 169 is connected to a scanning device 170 which is located in the switchcore 12 and is common to all switchports connected to matrix column 3. Accordingly, there is one scanning device for each of the columns of the switchcore matrix 12. The scanning device 170 scans the FIFO buffers 51 for any buffer that has a cell to be relayed to an output switchport 11b. The scanning device 170 reads the address contained in the header of the cells waiting in the buffers. Once the scanning device 170 finds one or more such buffers 51, it decides, using the priority rules discussed above in connection with FIG. 9, which buffer should be allowed to send its cell. That buffer 51 is pointed out by the output signal labeled "crosspoint select" in FIG. 23. The selected output switchport (SWP 10) will then start to extract the cell from the FIFO51.

[0104] The following logic controls access to the output switchport:

TABLE 1

POLL STATUS REGISTER	
SIGNAL	ACTION AT EVENT
output	Set RPF-bit output to: "1" if cell in crosspoint buffer 51 ("red light") "0" if no cell in crosspoint buffer 51 ("green light")
set	At the writing of the first byte of a cell into crosspoint buffer 51
R1 (if PRR 165 is "0")	At read-out of first byte of cell in crosspoint buffer 51
R2 (if PRR 165 is "1")	At read-out of last byte of cell in crosspoint buffer 51

TABLE 2

SCAN STATUS REGISTER	
SIGNAL	ACTION AT EVENT
output	Bit present in scan field: "1" - there is cell in crosspoint buffer 51 "0" - there is no cell in crosspoint buffer 51
S1 (if SRR 168 is "0")	At the writing of the first byte of a cell into crosspoint buffer 51
S2 (if SRR 168 is "1")	At the writing of the last byte of a cell into crosspoint buffer 51
Reset	At read of "first" byte of cell into crosspoint buffer 51

TABLE 3

SPEED COMBINATIONS		
	POLL RATE REGISTER	SCAN RATE REGISTER
i>o	"1" (R2)	"0" (S1)
i<o	"0" (R1)	"1" (S2)
i=o	0 (R1)	0 (S1)
i?o	1 (R2)	1 (S2)

[0105] In the above tables, positive logic has been used. From Table 1 it can be seen that as soon as a cell starts to be written into the buffer 51, the poll status register (PSR) 160 is set. As soon the first byte of a cell in the buffer 51

starts to be extracted, the PSR 160 is cleared if the poll rate register (PRR) 165 is at "0". If the PRR 165 is in position "1", the PSR 160 is cleared when the last byte is extracted from the buffer.

**[0106]** From Table 2 it can be seen that as soon as a cell starts to be written into the buffer 51 the scan status register (SSR) 162 is set at the first byte if the scan rate register (SRR) 168 is preloaded with "0". If the SRR 168 is preloaded with "1", the SSR 162 is set at the last byte of the cell in the crosspoint buffer 51. The SSR 162 is reset when the first byte starts to be read out from the buffer 51.

**[0107]** Table 3 is a state table indicating the states of the poll rate register (PRR) 165 and the scan rate register (SRR) 168 for different input and output bit rates. The contents of Table 3 are determined using the contents and definitions of Tables 1 and 2. If the input bit rate is larger than the output bit rate ( $i > o$ ) the following considerations apply: The readout of the cell in the FIFO buffer 51 can start immediately when the first byte of the cell arrives in the buffer. Therefore, the scan status register (SSR) 162 can be set immediately. From Table 2, at the line "At the writing of the first byte..." and from the corresponding column under SIGNAL, it follows that a "0" should be filled in under SCAN RATE REGISTER in Table 3 line  $i > o$ . A cell present in the FIFO buffer 51 is indicated by the fact that the poll rate register (PRR) 165 has the status of being occupied. Since the readout rate is slower than the input rate, the PRR 165 must not be cleared until the last byte has been read out from the buffer 51. From Table 1, last line, this gives the "1" state of the PRR 165, and this "1" is reflected in Table 3 under POLL RATE REGISTER, line  $i > o$ .

**[0108]** If, on the other hand, the PRR 165 was cleared when the first byte is extracted, the poll status register (PSR) 160 would receive an early indication that the FIFO buffer 51 is ready to receive a new cell. When this new cell is loaded into the buffer 51, there is a great risk of overwriting the old cell since the readout rate of the old cell is slower. This would, therefore, corrupt the data of the old cell. To prevent this, the PRR 165 is cleared at the last byte of the cell being extracted.

**[0109]** In the reversed situation, when the output bit rate "o" is higher than the input bit rate "i" ( $i < o$ ), the scan status register (SSR) 162 is not set until the last byte of the cell has been stored in the FIFO buffer 51. This means that there will be a "1" in the SSR 162 (next to last line in table 2) which gives the "1" in Table 3, second line, under SCAN RATE REGISTER. Given the fact that  $o > i$ , it is possible to clear the poll status register (PSR) 160 at readout of the first byte. When the PSR contents, i.e. the RPF-field 26, is analyzed in the switchport, the switchport is able to transfer a new cell into the buffer 51. The cell in the buffer 51 is read out before a new, slower, cell arrives in the buffer. From Table 1, next to last line "At read-out of first ..." and the corresponding entry under SIGNAL, a "0" is given and this "0" is reflected in Table 3 under POLL RATE REGISTER, line  $i < o$ .

**[0110]** If input and output rates are equal ( $i = o$ ), as soon as extraction of the first byte of the cell in the FIFO buffer 51 starts, the first byte of a new cell can be written into the buffer 51 with no risk of data corruption. Therefore, the value "0" under POLL RATE REGISTER, line  $i = o$  still applies. Thus, "0" is shown in line  $i = o$  below POLL RATE REGISTER. The scan rate register (SRR) 168 can be set at the arrival of the first byte of the cell in the buffer 51 giving the set condition S1, i.e., a "0" in Table 2. This "0" is entered below SCAN RATE REGISTER, line  $i = o$ , in Table 3.

**[0111]** When the input and output bit rates are unknown ( $i \neq o$ ), the scan status register (SSR) 162 cannot be set until the last byte of the cell has been written into the FIFO buffer 51. From Table 2, this gives the set condition S2, i.e., a "1". This "1" is entered in Table 3 under SCAN RATE REGISTER, line  $i \neq o$ . On the transmitting side, the poll status register (PSR) 160 cannot give permission for sending a new cell into the buffer 51 until the last byte of the cell being extracted has been read, as in reset condition R2 in Table 2. The "1" indicated at this condition is entered under POLL RATE REGISTER, line  $i \neq o$ , in Table 3. Accordingly, both the PSR 160 and the SSR 162 are set to "1" in this situation when  $i \neq o$ . This situation occurs at start up of the ATM switch and at the replacement of a device 17. In this way, the ATM switch configures and adapts itself to the bit rate of the replacement device 17. All other devices 17 remain unchanged and will operate at their individual bit rates.

**[0112]** In the situation where  $i = o$ , cells at the inlet switchport 11a can be concatenated or linked, even if they pass through the same matrix crosspoint, and the cells are switched through the ATM switch "on the fly". Concatenated cells have the head of a successive cell added to the tail of a previous cell.

**[0113]** Thus, where the present invention has been described in connection with the exemplary embodiments thereof, it can be understood that many modifications and variations will be apparent to those of ordinary skill in the art. The present disclosure and the following claims are intended to cover all such modifications and variations.

## Claims

### 1. A controlled access ATM switch comprising:

a switchcore matrix comprising a plurality of rows, columns, and crosspoints thereof, providing selectable routing paths for the routing of a stream of information cells from input points to output points on said matrix; a plurality of switchports electronically connected to said switchcore matrix at said input and output points, for

asynchronously transmitting and receiving said stream of information cells from said switchcore matrix, each of said switchports comprising:

5 means for interfacing between external information cell communication devices and said switch matrix; and  
means for controlling access to said switchcore matrix, said means for controlling access comprising:

one or more input buffers, selectively sized for the type of communication traffic to be transmitted through said buffer, for storing said information cells until selected routing paths in said switchcore matrix are free; characterized in

10 a relay address field for identifying target switchports to which an information cell is addressed;  
a relay poll field for identifying which of said target switchports are free;  
means for comparing said relay address field and said relay poll field to identify which of said target switchports are both addressed and free; and  
means for transmitting said information cell to said addressed target switchports that are free.

15 2. The controlled access ATM switch of claim 1 wherein the means for transmitting said information cells includes means for establishing priority among said information cells such that cells with the highest priority are transmitted first.

20 3. The controlled access ATM switch of claim 1 wherein the means for transmitting said information cells includes means for independently delaying transmission of selected ones of said information cells to synchronize said transmission with the transmission of other identified information cells to provide isochronal service.

25 4. The controlled access ATM switch of claim 1 wherein the means for transmitting said information cells includes a totally asynchronous interface between said switchports and said switchcore for enabling each of said external communications devices to operate at an optimum data rate.

5. The controlled access ATM switch of claim 1 wherein the means for transmitting said information cells includes:

30 means for phase shifting incoming and outgoing information cells; and  
means for determining the extent of said phase shift by measuring the length of time required to process and assemble said relay address field and said relay poll field.

35 6. The controlled access ATM switch of claim 5 wherein the switchcore matrix includes one or more single-cell buffers for each of said routing paths.

7. The controlled access ATM switch of claim 6 wherein said single-cell buffers are arranged in pools of common buffers at said input points of said switchcore matrix.

40 8. The controlled access ATM switch of claim 7 wherein each of said pools of common buffers comprise a group of between one and twelve of said single-cell buffers.

9. The controlled access ATM switch of claim 8 wherein each of said pools of common buffers are distributed across said switchcore matrix to the crosspoints that are used most often.

45 10. The controlled access ATM switch of claim 6 wherein each of said single-cell buffers is shared by two of said crosspoints in said switchcore matrix.

11. The controlled access ATM switch of claim 6 wherein one of said single-cell buffers is located at each of said crosspoints of said switchcore matrix.

50 12. The controlled access ATM switch of claim 6 wherein said switchcore matrix includes:

55 a plurality of row function units for terminating said information cell streams incoming from said switchports;  
a plurality of column function units which form synchronized pairs with said row function units and generate said information cell streams outgoing to said switchports;  
a plurality of cross function units for receiving said information cells from said row function units and relaying said information cells to said column function units; and

a plurality of row buses and column buses for transmitting said information cell streams between said row function units and said column function units.

5        13. The controlled access ATM switch of claim 12 wherein said cross function unit includes a two-port memory device for storing said information cell until said information cell is unloaded by said column function unit.

10      14. The controlled access ATM switch of claim 13 wherein said cross function unit includes means for generating a flag signal for indicating to said column function unit that an information cell addressed to said column function unit is being stored in said two-port memory device.

15      15. The controlled access ATM switch of claim 14 wherein said switchcore matrix is constructed on a single integrated circuit.

15      16. The controlled access ATM switch of claim 15 wherein said switchcore matrix is mounted on a backplane to which said switchports are electrically connected.

17. A method for controlling the flow of information cells within a communications system comprising the steps of:

20      providing selectable routing paths for the routing of said information cells from input points to output points of a switchcore matrix having a plurality of rows, columns, and crosspoints thereof;

20      electronically connecting a plurality of switchports to the input and output points of said switchcore matrix to transmit information cells thereto and receive information cells therefrom;

25      connecting each of said switchports to an external information cell communications device; and  
controlling access to said switchcore matrix available to each of said information cells, said step of controlling access comprising the steps of:

30      selectively sizing one or more input buffers located within each of said switchports for the type of communications traffic to be transmitted through said buffer;

30      storing of said information cells in said one or more input buffers located within each of said switchports until selected routing paths in said switchcore matrix are free; characterized in  
identifying target switchports to which an information cell is addressed;

35      identifying target switchports that are free;

35      comparing said addressed target switchports with said free target switchports to identify which of said addressed target switchports are free; and

35      transmitting said information cells to said addressed target switchports that are free.

35      18. The method as set forth in claim 17 wherein the step of transmitting said information cell includes:

40      measuring the length of time required to identify and compare said addressed target switchports with said free target switchports; and

40      phase shifting incoming and outgoing information cells an amount equivalent to said length of time required to identify and compare said addressed target switchports with said free target switchports.

45      19. The method as set forth in claim 18 wherein the step of providing selectable routing paths for the routing of information cells includes providing one or more single-cell buffers for each of said routing paths in said switchcore matrix.

50      20. The method as set forth in claim 19 further comprising the step of mounting said switchcore matrix, said single-cell buffers, and said double-directed connections on a single integrated circuit.

50      21. A data transfer switch comprising:

55      means for providing selectable routing paths for the routing of said information cells from input points to output points of a switchcore matrix having a plurality of rows, columns, and crosspoints thereof;

55      means for electronically connecting a plurality of switchports to the input and output points of said switchcore matrix to transmit information cells thereto and receive information cells therefrom;

55      means for connecting each of said switchports to an external information cell communications device; and  
means for controlling access to said switchcore matrix available to each of said information cells, said means of controlling access comprising:

means for selectively sizing one or more input buffers located within each of said switchports for the type of communications traffic to be transmitted through said buffer;  
means for storing said information cells in said one or more input buffers located within each of said switchports until selected routing paths in said switchcore matrix are free; characterized in  
means for identifying target switchports to which an information cell is addressed;  
means for identifying target switchports that are free;  
means for comparing said addressed target switchports with said free target switchports to identify which of said addressed target switchports are free; and  
means for transmitting said information cells to said addressed target switchports that are free.

10       **22.** The data switch of claim 21 wherein the means for transmitting said information cell includes:

15       means for measuring the length of time required to identify and compare said addressed target switchports with said free target switchports; and  
means for phase shifting incoming and outgoing information cells an amount equivalent to said length of time required to identify and compare said addressed target switchports with said free target switchports.

20       **23.** The data switch of claim 22 wherein the means for providing selectable routing paths for the routing of information cells includes means for providing one or more single-cell buffers for each of said routing paths in said switchcore matrix.

25       **24.** The data switch of claim 23 wherein the means for providing selectable routing paths for the routing of information cells includes means for constructing said switchcore matrix on a single integrated circuit.

30





35

40

45

50

55







FIG. 6




FIG. 7




FIG. 8




FIG. 8a





FIG. 8b

## Direction: Switch Port to Switch Core



## Switch Core to Switch Port





FIG. 8c

## Direction: Switch Port to Switch Core



## Switch Core to Switch Port









FIG. 8d









FIG.16




FIG.17





FIG.18



67

FIG. 19





FIG.20






FIG. 21 :




FIG. 22




FIG. 23

**THIS PAGE BLANK (USPTO)**



(19)

Europäisches Patentamt  
European Patent Office  
Office européen des brevets



(11)

EP 1 051 001 A3

(12)

## EUROPEAN PATENT APPLICATION

(88) Date of publication A3:  
07.03.2001 Bulletin 2001/10

(51) Int Cl.7: H04L 12/56, H04Q 11/04

(43) Date of publication A2:  
08.11.2000 Bulletin 2000/45

(21) Application number: 00111248.1

(22) Date of filing: 28.01.1994

(84) Designated Contracting States:  
AT BE CH DE DK ES FR GB GR IE IT LI LU MC NL  
PT SE

(30) Priority: 29.06.1993 US 11378

(62) Document number(s) of the earlier application(s) in  
accordance with Art. 76 EPC:  
94905904.2 / 0 681 770

(71) Applicant: TELEFONAKTIEBOLAGET LM  
ERICSSON  
126 25 Stockholm (SE)

(72) Inventor: Petersen, Lars-Göran  
147 42 Tumba (SE)

(74) Representative: Wennerholm, Kristian  
Ericsson Radio Systems AB  
Common Patent Department  
164 80 Stockholm (SE)

(54) Controlled access ATM switch

(57) An asynchronous transfer mode (ATM) switch in which access to a switchcore matrix is monitored and controlled through the logic and buffering functions of switchports connected thereto. The switchcore is greatly simplified by moving the logic and buffering functions to the switchports. The switchcore matrix comprises a plurality of rows, columns, and crosspoints thereof, pro-

viding routing paths for the routing of information cells from input points to output points on the matrix. Single-store buffers in the switchcore matrix enable temporary storage and hand-off of individual information cells as they pass through the matrix. The simplicity of the switchcore matrix enables it to be constructed on a single integrated circuit.




FIG. 3



European Patent  
Office

## EUROPEAN SEARCH REPORT

Application Number

EP 00 11 1248

DOCUMENTS CONSIDERED TO BE RELEVANT									
Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim	CLASSIFICATION OF THE APPLICATION (Int.Cl.)						
A	US 5 126 999 A (MUNTER ERNST A ET AL) 30 June 1992 (1992-06-30) * abstract * * column 3, line 36-61 *	1,17,21	H04L12/56 H04Q11/04						
A	US 4 692 917 A (FUJIOKA MASANOBU) 8 September 1987 (1987-09-08) * column 1, line 52 – column 2, line 16 *	1,17,21							
A	US 5 038 343 A (LEBIZAY GERALD ET AL) 6 August 1991 (1991-08-06) * abstract *	1,17,21							
A	US 5 079 762 A (TANABE SEN ICHI) 7 January 1992 (1992-01-07) * abstract *	1,17,21							
TECHNICAL FIELDS SEARCHED (Int.Cl.)									
H04L									
<p>The present search report has been drawn up for all claims</p> <table border="1" style="width: 100%; border-collapse: collapse;"> <tr> <td style="width: 33%;">Place of search:</td> <td style="width: 33%;">Date of completion of the search:</td> <td style="width: 33%;">Examiner:</td> </tr> <tr> <td>THE HAGUE</td> <td>16 January 2001</td> <td>Dhondt, E</td> </tr> </table> <p><b>CATEGORY OF CITED DOCUMENTS</b></p> <p>X : particularly relevant if taken alone  Y : particularly relevant if combined with another document of the same category  A : technological background  O : non-written disclosure  P : intermediate document</p> <p>T : theory or principle underlying the invention  E : earlier patent document, but published on, or after the filing date  D : document cited in the application  L : document cited for other reasons  B : member of the same patent family, corresponding document</p>				Place of search:	Date of completion of the search:	Examiner:	THE HAGUE	16 January 2001	Dhondt, E
Place of search:	Date of completion of the search:	Examiner:							
THE HAGUE	16 January 2001	Dhondt, E							

**ANNEX TO THE EUROPEAN SEARCH REPORT  
ON EUROPEAN PATENT APPLICATION NO.**

EP 00 11 1248

This annex lists the patent family members relating to the patent documents cited in the above-mentioned European search report. The members are as contained in the European Patent Office EDP file on The European Patent Office is in no way liable for these particulars which are merely given for the purpose of information.

16-01-2001

Patent document cited in search report		Publication date		Patent family member(s)	Publication date
US 5126999	A	30-06-1992	CA	1320257 A	13-07-1993
US 4692917	A	08-09-1987	JP	1899713 C	27-01-1995
			JP	6028361 B	13-04-1994
			JP	61127250 A	14-06-1986
			GB	2168222 A,B	11-06-1986
US 5038343	A	06-08-1991	EP	0405042 A	02-01-1991
			DE	68918275 D	20-10-1994
			JP	2699985 B	19-01-1998
			JP	3038137 A	19-02-1991
US 5079762	A	07-01-1992	JP	2531275 B	04-09-1996
			JP	3117137 A	17-05-1991

**THIS PAGE BLANK (USPTO)**